


# Integrated Computer Control System Startup Simulation

Brett M. Kettering

February 4, 1998



This is an informal report intended primarily for internal or limited external distribution. The opinions and conclusions stated are those of the author and may or may not be those of the Laboratory.

Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

#### DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This report has been reproduced  
directly from the best available copy.

Available to DOE and DOE contractors from the  
Office of Scientific and Technical Information  
P.O. Box 62, Oak Ridge, TN 37831  
Prices available from (423) 576-8401

Available to the public from the  
National Technical Information Service  
U.S. Department of Commerce  
5285 Port Royal Rd.,  
Springfield, VA 22161

---

# **Integrated Computer Control System Startup Simulation**

**FY98 LDRD Report**

---

**Brett M. Kettering  
February 4, 1998**

## Introduction

A simulation of the startup time for the ICCS (Integrated Computer Control System) applied to the National Ignition Facility (NIF) has been performed. The simulation results provide confidence that this requirement can be met. This document describes the results of the simulation that was done using the SIMPROCESS discrete-event-modeling tool.

The requirement for start up time established for the NIF control system is:

*Restarting all computers from a power-down condition shall take less than 30 minutes. This allows quick recovery from a power failure.*

## Overview

ICCS will run under both the Solaris and VxWorks operating systems. In most cases the operating system will be downloaded to a computer, rather than reside on local disks. Figure 1 graphically depicts the ICCS network modeled in the startup simulation.

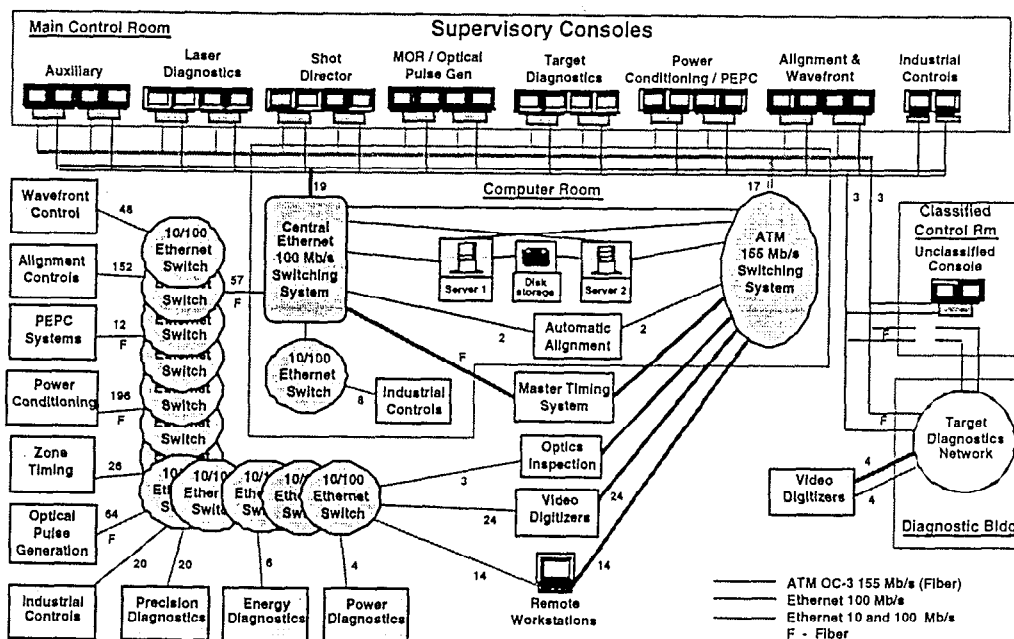


Figure 1 - ICCS Network Configuration

The model was run with several startup configurations. The results of the final configuration are reported in this document. Savings in startup time compared to other configurations are mentioned.

## Assumptions of the Model

Discrete-event simulations present some limitations in modeling complex, non-deterministic behaviors of time-sharing, multi-tasking computers and their operating systems. Some assumptions were made to simplify the model yet still maintain confidence that the result is reasonable. This section describes those assumptions.

### Computer System Hardware Availability

The simulation assumes that the computer system hardware is always operational. It does not interject hardware failures such as network switch breaking, computer will not boot, or mechanical disk failure. The MTBF (Mean Time Between Failure) figures published with this type of equipment is generally on the order of many years. It is our experience that failures of this kind are usually the result of human intervention.

### Distribution of Solaris

Computers referred to in this subsection can be seen in Figure 1. The two file servers boot from a local copy of the Solaris operating system and are two of the 63 Solaris computers in the network. The supervisory consoles, automatic alignment workstations, and several FEPs comprise the 61 Solaris computers that download the operating system from one of the four-CPU file servers. Thus, the 61 Solaris computers contend for a file server processor and access to the one copy of the operating system.

### Distribution of VxWorks FEP Images

The VxWorks FEPs do not have disks. They must download the operating system and applications from a Solaris computer on the network. Each Solaris computer will have its own local disk with a copy of VxWorks, NTP (Network Time Protocol) application, and all FEP applications. Thus when a VxWorks FEP boots and loads applications it does so from any one of the 63 Solaris computers.

### VxWorks FEP Download Resources Acquired for Duration of Download

A VxWorks FEP must acquire one of 63 Solaris resources to perform the information download and one of 57 network links to the main computer system location. The simulation assumes that these resources are acquired and held for the duration of the operating system download, then again for the NTP application download, and a third time for the FEP application download. Since each Solaris computer has an independent copy of this information on its own local disk when the Solaris computer resource is held the disk resource is also held.

The download operations use TFTP (Trivial File Transfer Protocol). This protocol actually downloads the information in packets which, in reality, allows multiple VxWorks FEPs to share the same network link and Solaris computer. This sharing is difficult to model and the simulation was not performed.

## Solaris Computers Boot First

All Solaris computers must complete booting before any VxWorks FEP may boot. This will yield a worst-case scenario.

In the actual system the Solaris computers will finish booting at different times. As the Solaris computers become available VxWorks FEPs will begin to boot. Thus, to some degree, Solaris computers and VxWorks FEPs will be booting in parallel.

## Supervisory Applications Run After VxWorks FEPs Initialize

All VxWorks FEPs must complete initialization before any supervisory application runs. This will yield a worst-case scenario.

In the actual system a supervisory application will run after the Solaris computer on which it resides completes booting and the VxWorks FEPs will become available at different times. The supervisory application will become ready when the VxWorks FEPs on which it depends become available. Thus, to some degree, the VxWorks FEPs and supervisory applications initialize in parallel.

## Derivation of Delays and Statistical Distributions

The simulation passes entities (e.g. booting a Solaris computer, requesting a CORBA reference, etc.) through a sequence of activities (e.g. retrieve an operating system from disk, download an application to a requester, etc.). Each activity has, among its other attributes, a delay. This delay defines how long it takes an entity to finish the activity. The sum of these delays from beginning to end is the total time required to complete the process modeled by the simulation. These delays can be constants or random numbers chosen from a statistical distribution.

This section describes how the constants and statistical distributions were derived for the startup simulation.

### Use of Beta Statistical Distribution

Startup activities have a minimum calculable duration and an expected maximum duration. The normal distribution allows duration that may exceed these bounds. The Beta distribution allows a minimum and maximum bound to be put on the distribution and a bell shape to be aligned where desired within those bounds. Thus, the Beta distribution was used to generate random delays in the startup simulation.

The Beta distributions were established based on measurements of similar activities and the experience of ICCS team members with computer systems.

### TFTP Download Times

A network performance tool called *Netperf* was used to measure *Request-Response* mode tests with 1 Kbyte packets. On a 10 Mbit/second Ethernet this test showed that 626 packets/second could be transferred. This test used TCP/IP (Transmission Control Protocol/Internet Protocol). TFTP uses UDP (User Datagram Protocol), which is a similar protocol, and 512 byte packets. It was assumed that TFTP could maintain the same packet transfer rate.

The VxWorks boot code was modified to time the boot parts of a VxWorks FEP. After subtracting an estimated disk access time during the TFTP download of the operating system, it was determined that the 626 packets/second transfer rate was approximately maintained, but that there was no difference in the rate for 10 Mbit/second and 100 Mbit/second networks.

### VxWorks FEP Boot Times

The VxWorks operating system is approximately 1 Mbyte. The modifications to the VxWorks boot code were used to determine the following:

- Time to initialize the FEP hardware before the operating system is downloaded
- Time to download the operating system
- Time to initialize the operating system after it is downloaded.

Two VxWorks FEPs were booted concurrently and no variance in the boot times was observed. The ICCS testbed does not currently have more than two VxWorks FEP on the same network.

Beta statistical distributions were developed for these parts based on several boot attempts. NTP is approximately the same size as the VxWorks operating system, so the same Beta statistical distributions for the download of the operating system were used for the download of NTP. It was estimated that a typical VxWorks FEP application would be approximately 8 Mbytes, based on a trivial sample VxWorks FEP application that is 6.6 Mbytes. A Beta distribution for the VxWorks FEP application TFTP download was created based on the information noted above.

### Solaris Computer Boot Times

The boot times for Solaris client computers were measured on ICCS computers that are similar to the ones expected in the system. The Solaris operating system image is approximately 10 Mbytes.

Beta distributions were created for disk access to read 10 Mbytes and for a TFTP download of 10 Mbytes. The remaining measured boot time was divided into two Beta distributions for hardware initialization and Solaris operating system initialization.

### Windows NT Computer Boot Times

The boot time for a 232 MHz Pentium Pro running Windows NT was measured. A Beta distribution based on those measurements was created for a standalone Windows NT computer.

### Startup and Initialization of Configuration Server

According to input from the FEP developers the simulation used 59,348 devices. The startup and initialization of the configuration server was measured with up to 40,000 devices. The times grew linearly with the number of devices. Thus, a Beta distribution was constructed using 1.5 multiplied by the time for 40,000 devices.

## FEP Application Initialization

At the time of the simulation, there was no CORBA product for VxWorks. Thus, the simulation used times measured with Solaris CORBA applications. The planned VxWorks FEPs will have faster processors and be on an equal bandwidth network as the Solaris applications that were measured.

Times for requesting, receiving, and initializing the controller and device objects were based upon the CORBA Fast-track project. The smallest delay unit of time in SIMPROCESS is a second. Since these times were in the millisecond range, constant values were used rather than statistical distributions.

## Shot Director and Supervisory Application Startup and Initialization

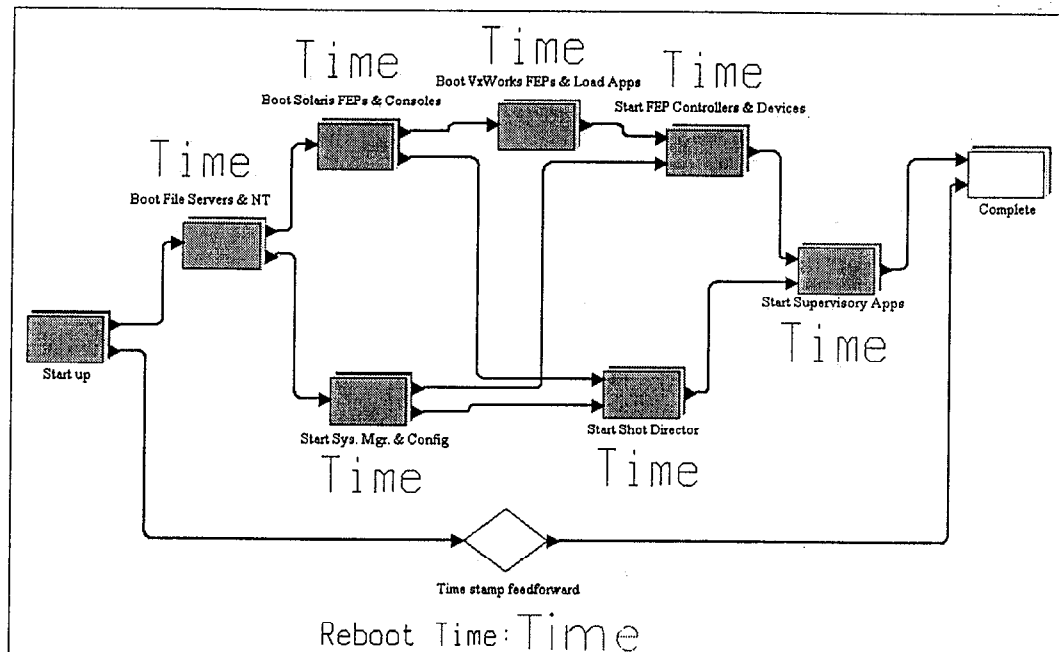
At the time of the simulation, no supervisory application has been prototyped. The numbers of objects that are created and initialized are best guesses.

Start times for the supervisory applications are Beta distributions based on measured start times of Solaris X/Motif applications. Times for requesting, receiving, and initializing the supervisory objects were based upon the CORBA Fast-track project. The smallest delay unit of time in SIMPROCESS is a second. Since these times were in the millisecond range, constant values were used rather than statistical distributions.

## Description of Startup Process

Figure 2 graphically depicts the startup simulation. Each rectangular box represents a SIMPROCESS *process*, which is a set of activities through which the startup entity must pass. The lines indicate *process* dependencies. This section describes each of these simulation processes.





**Figure 2 – Startup Simulation Flow of Processes**

### Start Up

This process generates the event that starts the simulation. One copy goes through the startup process and the other is forwarded along another path labeled *Time Stamp Feed Forward* to keep the time. The latter copy provides the total simulation time.

### Boot File Servers and NT

This process represents the computers that have their own copies of their operating systems and boot standalone. The file servers will have a large amount of memory and peripherals so that hardware checking and initialization will dominate their boot time.

### Boot Solaris FEPs and Consoles

This process represents booting all of the Solaris computers. They all boot in parallel and contend with one another for the file server and Solaris operating system disk resources.

### Start System Manager and Configuration Server

This process represents starting the initial file server programs. The system manager program runs first and starts the configuration server. The configuration server then reads the device configuration information out of the database and stores it in a memory-resident data structure.

## Boot VxWorks FEPs and Load Applications

This process represents booting the VxWorks FEPs and starting their applications. They contend with one another for the Solaris boot server and network resource to get their operating system and application. At the end of this process all the VxWorks FEPs applications are ready to be configured and initialized by the system manager and configuration server programs.

## Start Shot Director

This process represents starting the Shot Director application and creating its software objects. The Shot Director requires the system manager and configuration server programs to complete this process. The Shot Director application must be running before the supervisory applications are run because they depend on it to coordinate the shot process.

## Start FEP Controllers and Devices

This process represents configuring and initializing the software objects for both Solaris and VxWorks FEP applications. The FEP applications require the system manager and configuration server programs to complete this process. All FEP applications must be running before the supervisory applications that use them are run because the supervisory applications depend on the FEP applications to control the system components.

## Start Supervisory Applications


This process represents starting the supervisory applications and creating their software objects. They require the system manager and configuration server programs to complete this process.

## Complete

This process records the total simulation time and destroys the simulation entities, thus terminating one replication (the SIMPROCESS term for one run of the simulation) of the simulation.

## Results of the Simulation

Figure 3 graphically depicts the results of the startup simulation. Each time is given in minutes and is the total elapsed time when the process represented by the box completed. As indicated at the bottom, the total time to start up the system is 28.7 minutes for this replication. A series of 100 replications yielded a maximum of 30.0 minutes, minimum of 28.5 minutes, an average of 29.0 minutes, and a standard deviation of 0.5 minutes.



**Figure 3 – Startup Simulation with Time Indicators**

The simulation shows that, at most, having each Solaris computer with its own copy of the operating system saves 7.4 minutes in the system startup time. This would require an increase in the system administration effort to maintain 61 copies of the Solaris operating system. This saving is not incorporated in the configuration reported in this document.

It is understood that there is a system administration cost to maintaining 61 mirrored locations of the VxWorks FEP information. Nevertheless, there are 455 VxWorks FEPs and the simulation showed savings of 25.3 minutes in the system startup time over maintaining the information in a single location. This saving is incorporated in the configuration reported in this document.

## Conclusion

Under the assumptions and distributions of this simulation, the network and information distribution resources are sufficient to boot the computers and get their applications running within 30 minutes.

Generating more work in system administration may save some time. This decision must be based on how often an entire system startup is expected versus the cost and frequency of operating system maintenance and upgrade activities.

### **Appendix: Acronyms**

<b>Acronym</b>	<b>Definition</b>
FEP	Front End Processor
ICCS	Integrated Computer Control System
MTBF	Mean Time Between Failure
NIF	National Ignition Facility
NTP	Network Time Protocol
SSDR	Subsystem Design Requirement
TCP/IP	Transmission Control Protocol/Internet Protocol
TFTP	Trivial File Transfer Protocol
UDP	User Datagram Protocol